

101-790
A
09/60/98
U.S. PRO

Patent
Attorney's Docket No. 018414-229

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

**REQUEST FOR FILING CONTINUATION/DIVISIONAL
APPLICATION UNDER 37 C.F.R. § 1.53(b)**

Box PATENT APPLICATION

Assistant Commissioner for Patents
Washington, D.C. 20231

Sir:

This is a request for filing a [] continuation divisional application under 37 C.F.R. § 1.53(b) of pending Application No. 08/480,739 filed on June 7, 1995, for SYSTEM FOR PLACING ENTRIES OF AN OUTSTANDING PROCESSOR REQUEST INTO A FREE POOL AFTER THE REQUEST IS ACCEPTED BY A CORRESPONDING PERIPHERAL DEVICE, by the following named inventor(s):

- (a) Full Name William K. Zuravleff
- (b) Full Name Mark Semmelmeyer
- (c) Full Name Timothy Robinson
- (d) Full Name Scott Furman

- The entire disclosure of the prior application from which a copy of the oath or declaration is supplied herewith is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
- This application is being filed by less than all the inventors named in the prior application. In accordance with 37 C.F.R. 1.63(d)(2), the Commissioner is requested to delete the name(s) of the following person or persons who are not inventors of the invention being claimed in this application.

- (a) Full Name _____
- (b) Full Name _____
- (c) Full Name _____

1. Enclosed is a copy of the prior Application No. 08/480,739 as originally filed on June 7, 1995, including copies of the specification, claims, drawings and the executed oath or declaration as filed.
2. Enclosed is a revised prior application and a copy of the prior executed oath or declaration as filed. No new matter has been added to the revised application.

Request for Filing Continuation/Divisional Application
 of Application No. 08/480,739
 Attorney's Docket No. 018414-229
 Page 2

3. _____ statement(s) claiming small entity status are enclosed were filed in prior Application No. ___, filed on ___.
4. The filing fee is calculated below and in accordance with the enclosed preliminary amendment:

CLAIMS					
	NO. OF CLAIMS		EXTRA CLAIMS	RATE	FEE
Basic Application Fee					\$790.00
Total Claims	5	MINUS 20 =	0	x \$22.00 =	0.00
Independent Claims	1	MINUS 3 =	0	x \$82.00 =	0.00
If multiple dependent claims are presented, add \$270.00					0.00
Total Application Fee					\$790.00
If small entity status is claimed, subtract 50% of Total Application Fee					0.00
Add Assignment Recording Fee of \$40.00 if Assignment document is enclosed					0.00
TOTAL APPLICATION FEE DUE					\$790.00

5. Charge \$ 790.00 to Deposit Account No. 02-4800 for the fee due.
6. A check in the amount of \$ _____ is enclosed for the fee due.
7. The Commissioner is hereby authorized to charge any appropriate fees under 37 C.F.R. §§ 1.16, 1.17 and 1.21 that may be required by this paper, and to credit any overpayment, to Deposit Account No. 02-4800. This paper is submitted in triplicate.
8. Cancel in this application original claims 1-17 of the prior application before calculating the filing fee. (At least one original independent claim must be retained for filing purposes.)
9. Amend the specification by inserting before the first line the sentence: --This application is a continuation, divisional, of Application No. 08/480,739, filed June 7, 1995.--
10. Transfer the drawings from the pending prior application to this application and abandon said prior application as of the filing date accorded this application. A duplicate of this paper is

Request for Filing Continuation/Divisional Application
of Application No. 08/480,739
Attorney's Docket No. 018414-229
Page 3

enclosed for filing in the prior application file. (May only be used if signed by person authorized under 37 C.F.R. § 1.138 and before payment of issue fee.)

11. New drawings are enclosed.
12. Priority of Application No. __ filed on __ in __ (country) is claimed under 35 U.S.C. § 119.
 - The certified copy of the priority application
 is enclosed
 - was filed on __ in prior Application No. __, filed on __
 - has not yet been filed.
13. A preliminary amendment is enclosed.
14. Also enclosed _____
15. The power of attorney in the prior application is to Burns, Doane, Swecker & Mathis, L.L.P.
 - a. The power appears in the original papers in the prior application.
 - b. Since the power does not appear in the original papers, a copy of the power in the prior application is enclosed.
 - c. Recognize as Associate Attorney Robert S. Babayi, Reg. No. 33,471.

Request for Filing Continuation/Divisional Application
of Application No. 08/480,739
Attorney's Docket No. 018414-229
Page 4

d. Address all future communications to: (May only be completed by applicant, or attorney or agent of record.)

James W. Peterson
BURNS, DOANE, SWECKER & MATHIS, L.L.P.
P.O. Box 1404
Alexandria, Virginia 22313-1404

February 9, 1998
Date

By: 
James A. LaBarre
Registration No. 28,632

ADDRESS OF
SIGNATOR:

BURNS, DOANE, SWECKER & MATHIS, L.L.P.
P.O. Box 1404
Alexandria, Virginia 22313-1404
(703) 836-6620

inventor(s)
 assignee of complete interest
 attorney or agent of record
 filed under 37 C.F.R. § 1.34(a)

Patent
Attorney's Docket No. 018414-229

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of)
William K. ZURAVLEFF et al.) Group Art Unit: Unassigned
Application No.: Unassigned) Examiner: Unassigned
Filed: February 9, 1998)
For: SYSTEM FOR PLACING ENTRIES)
OF AN OUTSTANDING PROCESSOR)
REQUEST INTO A FREE POOL)
AFTER THE REQUEST IS ACCEPTED)
BY A CORRESPONDING)
PERIPHERAL DEVICE)

PRELIMINARY AMENDMENT

Assistant Commissioner for Patents
Washington, D.C. 20231

Sir:

Prior the examination of the above-identified U.S. patent application, please enter the following amendments.

IN THE TITLE:

Please change the title from "SYSTEM FOR PLACING ENTRIES OF AN OUTSTANDING PROCESSOR REQUEST INTO A FREE POOL AFTER THE REQUEST IS ACCEPTED BY A CORRESPONDING PERIPHERAL DEVICE" to --A METHOD FOR LIMITING ADMISSION OF PRIORITIZED MEMORY OR I/O TRANSACTIONS--.

IN THE CLAIMS:

Please cancel claims 1-17 without prejudice or disclaimer.

Please add the following new claims.

--21. A method according to Claim 18, further comprising the step of specifying a maximum number of outstanding memory or I/O transactions for said unique priority levels

in each of said queues which prevents entries of memory or I/O transactions having low priority levels from using one of said queues before entries of memory or I/O transactions having higher priority levels.

22. A method according to Claim 18, further comprising the step of determining priorities corresponding to entries of memory or I/O transactions by logical memory addresses, control bits derived from a memory management page table, control bits derived from segmentation entries, virtual addresses of a memory management system, programmable registers which set priorities for each processor, instructions, or instruction operands.--

REMARKS


Prior to the examination of the above-identified U.S. patent application, please enter the following preliminary amendments. The title of the application has been changed to better describe the application. Claims 21 and 22 have been added. The addition of these claims do not add new matter and is fully supported in the specification attached herewith.

Application No. Unassigned
Attorney's Docket No. 018414-229

Applicants believe that the application as filed is fully in condition for allowance, and early notification of such is hereby earnestly requested.

Respectfully submitted,

BURNS, DOANE, SWECKER & MATHIS, L.L.P.

By: 

Robert S. Babayi
Registration No. 33,471

P.O. Box 1404
Alexandria, Virginia 22313-1404
(703) 836-6620

Date: February 9, 1998

**COMBINED DECLARATION AND POWER OF ATTORNEY
FOR UTILITY PATENT APPLICATION**

Attorney Docket No.

018414-148

As a below-named inventor, I hereby declare that:
My residence, post office address and citizenship are as stated below next to my name;
I BELIEVE I AM THE ORIGINAL, FIRST AND SOLE INVENTOR (if only one name is listed below) OR AN
ORIGINAL, FIRST AND JOINT INVENTOR (if more than one name is listed below) OF THE SUBJECT MATTER
WHICH IS CLAIMED AND FOR WHICH A PATENT IS SOUGHT ON THE INVENTION ENTITLED:

NON-BLOCKING LOAD BUFFER AND A MULTIPLE-PRIORITY MEMORY SYSTEM

FOR REAL-TIME MULTIPROCESSING

the specification of which

(check one) is attached hereto;
 was filed on _____ as
Application No. _____
and was amended on _____;
(if applicable)

I HAVE REVIEWED AND UNDERSTAND THE CONTENTS OF THE ABOVE-IDENTIFIED SPECIFICATION,
INCLUDING THE CLAIMS, AS AMENDED BY ANY AMENDMENT REFERRED TO ABOVE;

I ACKNOWLEDGE THE DUTY TO DISCLOSE TO THE OFFICE ALL INFORMATION KNOWN TO ME TO BE
MATERIAL TO PATENTABILITY AS DEFINED IN TITLE 37, CODE OF FEDERAL REGULATIONS, Sec. 1.56
(as amended effective March 16, 1992);

I do not know and do not believe the said invention was ever known or used in the United States of America before
my or our invention thereof, or patented or described in any printed publication in any country before my or our
invention thereof or more than one year prior to said application; that said invention was not in public use or on sale in
the United States of America more than one year prior to said application; that said invention has not been patented or
made the subject of an inventor's certificate issued before the date of said application in any country foreign to the
United States of America on any application filed by me or my legal representatives or assigns more than twelve
months prior to said application;

I hereby claim foreign priority benefits under Title 35, United States Code Sec. 119 and/or Sec. 365 of any foreign
application(s) for patent or inventor's certificate as indicated below and have also identified below any foreign
application for patent or inventor's certificate on this invention having a filing date before that of the application(s) on
which priority is claimed:

COMBINED DECLARATION AND POWER OF ATTORNEY

Attorney Docket No.

018414-148

COUNTRY/INTERNATIONAL	APPLICATION NUMBER	DATE OF FILING (day, month, year)	PRIORITY CLAIMED
			YES <u> </u> NO <u> </u>
			YES <u> </u> NO <u> </u>

I hereby appoint the following attorneys and agent(s) to prosecute said application and to transact all business in the Patent and Trademark Office connected therewith and to file, prosecute and to transact all business in connection with international applications directed to said invention:

William L. Mathis	17,337	Regis E. Sluter	26,999	James W. Peterson	26,057
Peter H. Smolka	15,913	Samuel C. Miller, III	27,360	Teresa Stanek Rea	30,427
Robert S. Swecker	19,885	Ralph L. Freedland, Jr.	16,110	Robert E. Krebs	25,885
Platon N. Mardros	22,124	Robert G. Mukai	28,531	Robert M. Schulman	31,196
Benton S. Duffee, Jr.	22,030	George A. Hovance, Jr.	28,223	William C. Rowland	30,888
Joseph R. Magnone	24,239	James A. LaBarre	28,632	T. Gene Dillahunt	25,423
Norman H. Steno	22,716	E. Joseph Gess	28,510	Anthony W. Shaw	30,104
Ronald L. Grudziecki	24,970	David D. Reynolds	29,273	Patrick C. Keane	32,858
Frederick G. Michaud, Jr.	26,003	R. Danny Huntington	27,903	Bruce J. Boggs, Jr.	32,344
Alan E. Kopecki	25,813	Eric H. Weisblatt	30,505	William H. Benz	25,952

and: Mark Superko, Registration No. 34,027

Address all correspondence to: James A. LaBarre
Burns, Doane, Swecker & Mathis
P.O. Box 1404
Alexandria, Virginia 22313-1404

Address all telephone calls to: James A. LaBarre at (703) 836-6620.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

FULL NAME OF SOLE OR FIRST INVENTOR	SIGNATURE <u>William K. Zuravleff</u>	DATE <u>6/15/95</u>
William K. Zuravleff	CITIZENSHIP	
RESIDENCE	United States	
Mountainview, California		
POST OFFICE ADDRESS		
89 Mercy Street, Mountainview, California 94041		
FULL NAME OF SECOND JOINT INVENTOR, IF ANY	SIGNATURE <u>Mark Semmelmeyer</u>	DATE <u>08 June 95</u>
Mark Semmelmeyer	CITIZENSHIP	
RESIDENCE	United States	
Sunnyvale, California 94086		
POST OFFICE ADDRESS		
728 Silver Pine Court, Sunnyvale, California 94086		
FULL NAME OF THIRD JOINT INVENTOR, IF ANY	SIGNATURE <u>Timothy B. Roberts</u>	DATE <u>7/2/95</u>
Timothy B. Roberts	CITIZENSHIP	
RESIDENCE	United States	
Boulder Creek, California 95006		
POST OFFICE ADDRESS		
16 Blackstone Drive, Boulder Creek, California 95006		

FULL NAME OF FOURTH JOINT INVENTOR, IF ANY Scott FURMAN RESIDENCE Union City, California 94587		SIGNATURE  CITIZENSHIP United States	DATE 6/1/95
POST OFFICE ADDRESS 33 Union Street, #1412, Union City, California 94587			
FULL NAME OF FIFTH JOINT INVENTOR, IF ANY		SIGNATURE	
RESIDENCE		CITIZENSHIP	
POST OFFICE ADDRESS			
FULL NAME OF SIXTH JOINT INVENTOR, IF ANY		SIGNATURE	
RESIDENCE		CITIZENSHIP	
POST OFFICE ADDRESS			
FULL NAME OF SEVENTH JOINT INVENTOR, IF ANY		SIGNATURE	
RESIDENCE		CITIZENSHIP	
POST OFFICE ADDRESS			
FULL NAME OF EIGHTH JOINT INVENTOR, IF ANY		SIGNATURE	
RESIDENCE		CITIZENSHIP	
POST OFFICE ADDRESS			
FULL NAME OF NINTH JOINT INVENTOR, IF ANY		SIGNATURE	
RESIDENCE		CITIZENSHIP	
POST OFFICE ADDRESS			
FULL NAME OF TENTH JOINT INVENTOR, IF ANY		SIGNATURE	
RESIDENCE		CITIZENSHIP	
POST OFFICE ADDRESS			
FULL NAME OF ELEVENTH JOINT INVENTOR, IF ANY		SIGNATURE	
RESIDENCE		CITIZENSHIP	
POST OFFICE ADDRESS			
FULL NAME OF TWELFTH JOINT INVENTOR, IF ANY		SIGNATURE	
RESIDENCE		CITIZENSHIP	
POST OFFICE ADDRESS			

NON-BLOCKING LOAD BUFFER AND A MULTIPLE-PRIORITY MEMORY SYSTEM FOR REAL-TIME MULTIPROCESSING

FIELD OF THE INVENTION

The present invention relates to interfacing between a microprocessor, 5 peripherals and/or memories. More particularly, the present invention relates to a data processing system for interfacing a high-speed data bus, which is connected to one or more processors and a cache, to a second, possibly lower speed peripheral bus, which is connected to a plurality of peripherals and memories, by a non-blocking load buffer so that a continuous maximum throughput from multiple I/O 10 devices is provided via programmed I/O (memory-mapped I/O) of the processors and a plurality of non-blocking loads may be simultaneously performed. Also, the present invention is directed to a multiple-priority non-blocking load buffer which serves a multiprocessor for running real-time processes of varying deadlines by priority-based non-FIFO scheduling of memory and peripheral accesses.

BACKGROUND OF THE INVENTION

Conventionally, microprocessor and memory system applications retrieve data from relatively low bandwidth I/O devices, process the data by a processor and then store the processed data to the low bandwidth I/O devices. In typical 20 microprocessor memory system applications, a processor and a cache are directly coupled by a bus to a plurality of peripherals such as I/O devices and memory devices. However, the processor may be unnecessarily idled due to the latency between the I/O devices and the processor, thus causing the processor to stall and, as a result, excessive time is required to complete tasks.

In known processing systems, when an operation is performed on one of the 25 peripherals such as a memory device, the time between performing such an operation and subsequent operations is dependent upon the latency period of the memory device. Thereby, the processor will be stalled during the entire duration of the memory transaction. One solution for improving the processing speed of known

processing systems is to perform additional operations during the time of the latency period as long as there is no load-use dependency upon the additional operations. For example, if data is loaded into a first register by a first operation (1), where the first operation (1) corresponds to:

5 load r1<-[r2] (1)

and the first register is added to another operation by a second operation (2) where the second operation (2) corresponds to:

add r3←r1+r4 (2)

10 the operation (2) is load-use dependent and the operation (2) must wait for the latency period before being performed. Figures 1(a) and 1(b) illustrate a load-use dependent operation where the time for initiating the operation (2) must wait until the operation (1) is completed. Operation (2) is dependent upon the short latency period t_s in Figure 1(a) corresponding to a fast memory and the long latency period t_l in Figure 1(b) corresponds to a slower memory.

15 A non-blocking cache and a non-blocking processor are known where a load operation is performed and additional operations other than loads may be subsequently performed during the latency period as long as the operation is not dependent upon the initial load operation. Figures 2(a) and 2(b) illustrate such operations. In operation (1), the first register is loaded. Next, operations (1.1) and (1.2) are to be executed. As long as operations (1.1) and (1.2) are not load dependent on another load, these operations may be performed during the latency period t_2 as illustrated in Figure 2(a). However, if operation (1.1) is either load-dependent or a pending load, operations (1.1) and (1.2) must wait until the latency period t_1 ends before being performed.

25 Also known is a Stall-On-Use (Hit Under Miss) operation for achieving cache miss optimizations as described in "A 200 MFLOP Precision Architecture Processor" at Hot Chips IV, 1993, William Jaffe et al. In this Hit Under Miss operation, when one miss is outstanding only certain other types of instructions may be executed before the system stalls. For example, during the handling of a load

miss, execution proceeds until the target register is needed as an operand for another instruction or until another load miss occurs. However, this system is not capable of handling two misses being outstanding at the same time. For a store miss, execution proceeds until a load or sub-word store occurs to the missing line.

5 This Hit Under Miss feature can improve the runtime performance of general-purpose computing applications. Examples of programs that benefit from the Hit Under Miss feature are SPEC benchmarks, SPICE circuit simulators and gcc C compilers. However, the Hit Under Miss feature does not sufficiently meet the high I/O bandwidth requirements for digital signal processing applications such as 10 digital video, audio and RF processing.

Known microprocessor and memory system applications use real-time processes which are programs having deadlines corresponding to times where data processing must be completed. For example, an audio waveform device driver process must supply audio samples at regular intervals to the output buffers of the 15 audio device. When the driver software is late in delivering data for an audio waveform device driver, the generated audio may be interrupted by objectionable noises due to an output buffer underflowing.

In order to analyze whether or not a real-time process can meet its deadlines under all conditions requires predictability of the worst-case performance of the real-time processing program. However, the sensitivity of the real-time processing 20 program to its input data or its environment makes it impractical in many cases to exhaustively check the behavior of the process under all conditions. Therefore, the programmer must rely on some combination of analysis and empirical tests to verify that the real-time process will complete in the requisite time. The goals of real-time processing tend to be incompatible with computing platforms that have memory or 25 peripheral systems in which the latency of the transactions is unpredictable because an analysis of whether the real-time deadlines can be met may not be possible or worst-case assumptions of memory performance are required. For example, performance estimates can be made by assuming that every memory transaction

takes the maximum possible time. However, such an assumption may be so pessimistic that any useful estimate for the upper bound on the execution time of a real-time task cannot be made. Furthermore, even if the estimates are only slightly pessimistic, overly conservative decisions will be made for the hardware 5 performance requirements so that a system results that is more expensive than necessary.

Also, it is especially difficult to reliably predict real-time processing performance on known multiprocessors because the memory and peripherals are not typically multi-ported. Therefore simultaneous access by two or more processors 10 to the same memory device must be serialized. Even if a device is capable of handling multiple transactions in parallel, the bus shared by all of the processors may still serialize the transactions to some degree.

If memory requests are handled in a FIFO manner by a known multiprocessor, a memory transaction which arrives slightly later than another 15 memory transaction may take a much longer amount of time to complete since the later arriving memory requests must wait until the earlier memory request is serviced. Due to this sensitivity, very small changes in the memory access patterns of a program can cause large changes in its performance. This situation grows worse as more processors share the same memory. For example, if ten processors 20 attempt to access the same remote memory location simultaneously, the spread in memory latency among the processors might be 10:1 because as many as nine of these memory transactions may be buffered for later handling. In general, it is not possible to predict which of these processors will suffer the higher latencies and which of these processors will receive fast replies to their memory accesses. Very 25 small changes to a program or its input data may cause the program to exhibit slight operation differences which perturb the timing of the memory transactions.

Furthermore, types of memory which exhibit locality effects may exacerbate the above-described situation. For example, accesses to DRAMs are approximately two times faster if executed in page mode. To use page mode, a recent access must

have been made to an address in the same memory segment (page). One of the most common access patterns is sequential accesses to consecutive locations in memory. These memory patterns tend to achieve high page locality, thus achieving high throughput and low latency. Known programs which attempt to take advantage 5 of the benefits of page mode may be thwarted when a second program executing on another processor is allowed to interpose memory transactions on a different memory page. For instance, if ten processors, each with its own sequential memory access pattern, attempt to access the same DRAM bank simultaneously and each of the accesses is to a different memory page, the spread and memory latencies 10 between the fastest and slowest responses might be more than 25:1.

The present invention is directed to allowing a high rate of transfer to memory and I/O devices for tasks which have real-time requirements. The present invention is also directed to allowing the system to buffer I/O requests from several processors within a multiprocessor at once with a non-blocking load buffer. 15 Furthermore, the present invention is directed to extending the basic non-blocking load buffer to service a data processing system running real-time processes of varying deadlines by using scheduling of memory and peripheral accesses which is not strictly FIFO scheduling.

SUMMARY OF THE INVENTION

20 An object of the present invention is to reduce the effect of memory latency by overlapping a plurality of non-blocking loads during the latency period and to achieve a similar reduction on the effect of latency for non-blocking stores without requiring the additional temporary storage memory of a separate store buffer.

25 Another object of the present invention is to provide a non-blocking load buffer which buffers I/O requests from one or more processors within a multiprocessor so that a plurality of I/O memory transactions, such as loads and stores, may be simultaneously performed and the high I/O bandwidth requirements for digital signal processing applications may be met.

A still further object of the present invention is to provide a multiple-priority non-blocking load buffer for serving a multiprocessor running real-time processes of varying deadlines by using a priority-based method to schedule memory and peripheral accesses.

5 The objects of the present invention are fulfilled by providing a data processing system comprising a first data bus for transferring data requests at a first speed, a second bus for transmitting I/O data of a second speed, and a non-blocking load buffer connected to the first and second data buses for holding the data requests and the I/O data so that a plurality of loads and stores may be performed simultaneously. It is possible for the speed of the second bus to be slower than the first speed of the first bus. As a result, data may be retrieved from relatively low bandwidth I/O devices and the retrieved data may be processed and stored to low bandwidth I/O devices without idling the system unnecessarily.

10 In a further embodiment for a data processing system of the present invention, the first data bus is connected to a processor and the second data bus is connected to a plurality of peripherals and memories. The data processing system for this embodiment reduces the effect of latency of the peripherals and memories so that the application code, which uses programmed I/O, may meet its real-time constraints. Furthermore, the code may have reduced instruction scheduling constraints and a near maximum throughput may be achieved from the I/O devices of varying speeds.

15 Another embodiment of the present invention is fulfilled by providing a non-blocking load buffer comprising a memory array for temporarily storing data, and a control block for simultaneously performing a plurality of data loads and stores. 20 The non-blocking load buffer for this embodiment allows I/O requests from several processors to be performed at once.

25 Further scope of applicability of the present invention will become apparent from the detailed description given hereafter. However, it should be understood that the detailed description and specific examples, while indicating preferred

embodiments of the invention, are given by way of illustration only, since various changes and modifications within the spirit and scope of the invention will become apparent to those skilled in the art from this detailed description.

BRIEF DESCRIPTION OF THE DRAWINGS

5 The present invention will become more fully understood from the detailed description given hereinbelow and the accompanying drawings which are given by way of illustration only, and thus are not limitative of the present invention and wherein:

10 Figures 1(a) and 1(b) illustrate time dependency in a conventional load use operation;

Figures 2(a) and 2(b) illustrate time dependency in a known Stall-On-Use operation;

Figure 2(c) illustrates time dependency in a non-blocking load buffer for an embodiment of the present invention;

15 Figure 3(a) illustrates unpredictable memory latency in a known multiprocessor system;

Figure 3(b) illustrates memory latency in a multiprocessor system having a multiple-priority version of the non-blocking load buffer;

20 Figure 4 illustrates a block diagram for the data processing system according to one embodiment of the present invention;

Figure 5 is a schematic illustration of the non-blocking load buffer for an embodiment of the present invention;

Figure 6 illustrates an example of the contents for the entries in the memory of the data processing system for an embodiment of the present invention;

25 Figures 7(a) and 7(b) illustrate examples of the possible states for a non-blocking load through the memory of the data processing system for an embodiment of the present invention;

Figures 7(c), 7(d), 7(e), 7(f), and 7(g) illustrate the progress of addresses through queues of the data processing system in an embodiment of the present invention;

5 Figure 8 illustrates the circuitry used for the non-blocking load buffer in an embodiment of the present invention;

Figures 9(a) and 9(b) illustrate parallel pending queues which allow prioritization of data in the data processing system for an embodiment of the present invention;

10 Figure 10 illustrates a detailed block diagram of an embodiment of the non-blocking load buffer; and

Figure 11 illustrates a flow chart of the control of the pending queue in an embodiment of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

15 The first embodiment of the present invention will be discussed with reference to Figure 4. The data processing system includes a non-blocking load buffer 10, a cache 20, one or more processors 30₀, 30₁, ..., 30_m, a processor/cache bus 40 for connecting the non-blocking load buffer 10, the cache 20 and the processors 30_{0..m}, a plurality of peripheral devices 50₀, 50₁, ..., 50_n, and a peripheral bus 60, which includes an output bus 61 and an input bus 62, for connecting the 20 non-blocking load buffer 10 with the plurality of peripheral devices 50_{0..n}. Some or all of the peripheral devices 50_{0..n} may be memories, such as DRAM for example, which serve as a backing store for the cache 20 or as memory which the processors 30_{0..m} access directly through the non-blocking load buffer 10, thus bypassing the cache 20.

25 The processor/cache bus 40, which includes an input bus 41 and an output bus 42, transfers data from the cache 20 to registers of the processors 30_{0..m} as loads via the output bus 42 or alternatively, from the registers of the processors 30_{0..m} to the cache 20 as stores via the input bus 41. All I/O operations, which are programmed I/O, and any memory operations, even those which bypass the cache

20, are transmitted over the processor/cache bus 40 as well. For such operations the source or sink of the data is one of the peripheral devices 50_{0..n}, such as a DRAM, for example.

5 The bandwidth of the peripheral bus 60 should be chosen to meet the peak input or output bandwidth which is to be expected by the peripheral devices 50_{0..n}. The bandwidth of the non-blocking load buffer 10 should therefore be chosen to at least equal the sum of the bandwidth of the processor/cache bus 40 and the bandwidth of the peripheral bus 60.

10 The processors 30_{0..m} may be designed with dependency-checking logic so that they will only stall when data requested by a previously executed load instruction has not yet arrived and that data is required as a source operand in a subsequent instruction. Figure 2(c) illustrates the time dependency for the non-blocking load buffer 10 in an embodiment of the present invention. In Figure 2(c), load instructions are executed at operations (1), (1.1), (1.2), ... (1.n). In this 15 example, operations (1), (1.1), (1.2), ... (1.n) correspond to:

```
load r1<-[r2] (1);
load r5<-[r2 + 24] (1.1);
load r6<-[r2 + 32] (1.2);
. . . function z1 (1.n)
```

20 In this example, operation (2) executes an instruction dependent on operation (1), operation (2.1) executes an instruction dependent on operation (1.1), operation (2.2) executes an instruction dependent on operation (1.2), ... operation (2.n) executes an instruction dependent on operation (1.n). For instance, operations (2.1), (2.2), ... (2.n) correspond to:

25

```
add r3 <- r1 + r4 (2);
add r7 <- r5 + r4 (2.1);
add r8 <- r6 + r4 (2.2);
. . . function z2 (2.n).
```

n may be any positive integer. Operations (1.1), (1.2), ... (1. n) occur as fast as possible and the time between these operations is dependent upon the processor cycle for each operation. As shown in Figure 2(c), operation (2) waits only for a part of the latency period t_2 between operations (1) and (2), which is represented by 5 t_i , since the part of the latency period t_2 between operations (1) and (2) is overlapped with the operations (1.1), (1.2), ... (1. n). Thereafter, operations (2.1) and (2.2) occur as fast as possible after operation (2) with the time limitation being the processor cycle. The operations (2.1) and (2.2) do not wait for the full latency period t_2 of operation (1.1) since the latency period t_2 is partially overlapped with 10 operations (1.1), (1.2), ... (1. n) and the latency period t_2 for operation (1). For values of n sufficiently large, t_i may effectively be reduced to zero and the processor will not suffer any stalls due to memory or I/O latency. In this embodiment, there may be more than one outstanding load to slower memories or peripherals, as 15 shown at instant t_x in Figure 2(c) for example, and requests may thereby be pipelined to provide continuous maximum throughput from the plurality of peripheral devices $50_{0..n}$ from programmed I/O of the processor.

The non-blocking load buffer 10 may accept both load and store requests from the processors $30_{0..m}$ to the peripheral devices $50_{0..n}$. The non-blocking load buffer 10 may accept requests at the maximum rate issued by a processor for a 20 period of time that is limited by the amount of intrinsic storage in the non-blocking load buffer 10. After accepting requests, the non-blocking load buffer 10 then forwards the requests to the peripheral devices $50_{0..n}$ at a rate which the peripheral devices $50_{0..n}$ may process them. However, the non-blocking load buffer 10 does not have to wait until all of the requests are accepted.

25 Figure 5 illustrates internal architecture for the non-blocking load buffer in an embodiment of the present invention. The internal architecture of the non-blocking load buffer 10 includes a plurality of variable-length pending queues 114_0 , 114_1 , ... 114_n which contain entries that represent the state of requests made by the processors $30_{0..m}$ to the peripheral devices $50_{0..n}$. The pending queues $114_{0..n}$

correspond to each of the peripheral devices $50_{0..n}$ and are used to hold the address and control information for a load or a store before being transmitted to the corresponding peripheral device. If an entry in the pending queues $114_{0..n}$ contains a store transaction, then that entry will also contain the data to be written to the 5 corresponding peripheral device. Another component of the non-blocking load buffer 10 is a variable-length return queue 116 used to buffer the data returned by any of the peripheral devices $50_{0..n}$ in response to a load until the data can be transmitted to the requesting processor. In addition, there is a free pool 112 that contains any unused request entries. Once a load request is sent from one of the 10 pending queues $114_{0..n}$ to its corresponding peripheral device, the load request is marked as "outstanding" until a return data value is sent by the peripheral and enqueued in the return queue 116. At any time, a given request entry is either outstanding, in one of the pending queues $114_{0..n}$, in the return queue 116, or is not in use and is therefore located in the free pool 112.

15 Figure 6 illustrates an example of a memory structure 200 for storing the information present in each entry of the non-blocking load buffer 10. E entries may be stored in the non-blocking load buffer 10 and each entry holds one of a plurality of requests $201_1, 201_2, \dots 201_E$ which can either be a load or a store to any of the peripheral devices $50_{0..n}$. Each request 201 includes a control information 202, an 20 address 204 and data 206. However, for a load operation, the data 206 is empty until data returns from the peripheral devices $50_{0..n}$. The control information 202 describes parameters such as whether the transaction is a load or a store, the size of data transferred, and a coherency flag that indicates whether all previous transactions must have been satisfied before the new transaction is permitted to be 25 forwarded to the addressed peripheral device. The coherency flag enforces sequential consistency, which is occasionally a requirement for some I/O or memory operations. The data stored in the entries 201 need not be transferred among the queues of the non-blocking load buffer 10. Rather pointers to individual entries

within the memory structure 200 may be exchanged between the pending queues 114_{0...n}, the return queue 116 and the free pool 112 for example.

Figures 7(a) and 7(b) illustrate examples of the operation of the non-blocking load buffer 10 which will be input to a peripheral device such as a RAM. Initially, 5 all entries of the non-blocking load buffer 10 are unused and free. With reference to Figure 7(a), when one of the processors 30_{0...m} issues a non-blocking load, a free entry is chosen and the address of the load is stored in this entry. Additionally, space for the returning data is allocated. This entry now enters the pending queue, one of the queues 114_{0...n}, corresponding to the addressed peripheral device. When 10 the peripheral device is ready, the pending entry is delivered over the peripheral bus 60 and becomes an outstanding entry. When the outstanding entry returns from the peripheral device via the peripheral bus 60, the accompanying requested load data is written to the part of the entry that had been allocated for the returning data and the entry is placed in the return queue 116. The data is buffered until the one 15 processor is ready to receive the data over the processor/cache bus 40 (i.e., until the one processor is not using any required cache or register file write ports). Finally, the data is read out and returned to the one processor over the processor/cache bus 40 and the entry again becomes free. In addition, a peripheral device may accept multiple requests before returning a response.

20 Figure 7(b) illustrates the progress of an entry during a store operation. The store operation is similar to the load operation but differs because data does not need to be returned to the processors 30_{0...m} and this return step is, therefore, skipped. Once a store is accepted by the non-blocking load buffer 10, the store is treated by the issuing processor as if the data has already been transferred to the target 25 peripheral device, one of the peripheral devices 50_{0...n}.

The non-blocking load buffer 10 performs similarly to a rate-matching FIFO but differs in several ways. For instance, one difference is that although memory transactions forwarded by the non-blocking load buffer 10 are issued to any particular peripheral device 50_{0...n} in the same order that they were requested by any

of the processors $30_{0..m}$, the transactions are not necessarily FIFO among different peripheral devices. This is similar to having a separate rate-matching FIFO for outputting to each of the peripheral devices $50_{0..n}$ except that the storage memory for all the FIFOs is shared and is thereby more efficient. Also, the peripheral devices 5 $50_{0..n}$ are not required to return the results of a load request to the non-blocking load buffer 10 in the same order that the requests were transmitted.

The non-blocking load buffer 10 is capable of keeping all of the peripheral devices $50_{0..n}$ busy because an independent pending queue exists for each of the peripheral devices. A separate flow control signal from each of the peripheral devices indicates that this corresponding peripheral device is ready. Accordingly, 10 a request for a fast peripheral device does not wait behind a request for a slow peripheral device.

Figures 7(c), 7(d), 7(e), 7(f) and 7(g) illustrate examples of a non-blocking load progressing through the queues. After the queues are reset, entries 1, 2, 3 and 15 4 are in the free pool as shown in Figures 7(c). A non-blocking load is pending after the entry 1 is written into one of the pending queues as shown in Figure 7(d). Next, Figure 7(e) illustrates that entry 1 becomes outstanding after the pending load is accepted by its corresponding peripheral device. Subsequently, the load is returned from the peripheral device and entry 1 is now in the returned queue as 20 illustrated in Figure 7(f), and awaits transmission to the requesting processor (or to the cache if the access was part of a cache fill request). Finally, in Figure 7(g), entry 1 is released to the free pool and, once again, all entries are free.

Figure 8 illustrates the circuitry used for the non-blocking load buffer in an embodiment of the present invention which implements variable depth FIFOs. This 25 circuitry manipulates pointers (i.e. RAM addresses for example) to entries in a RAM 130 instead of storing the addresses of only the head and tail of a contiguous list of entries in the RAM 130. Furthermore, the circuitry does not require the pointers to be sequential. Therefore, records need not be contiguous and the RAM 130 can contain several interleaved data structures allowing more flexible allocation

and de-allocation of entries within a FIFO so that it is possible to use the same storage RAM 130 for multiple FIFOs. The entire circuit in Figure 8 represents a large FIFO of an unspecified data width which is controlled by a smaller FIFO represented by block 170 having pointer widths smaller than the records. In Figure 5 8, a first control circuit 110, a second control circuit 120, the RAM 130 having separate read and write ports, a plurality of shift registers 150₁, 150₂, 150₃, ...150_n corresponding to each entry in the RAM 130 and a multiplexer 160 are illustrated. The control circuits 110 and 120 provide pointers to the entries of the RAM 130 rather than the entries themselves being queued. Block 170 is essentially a FIFO 10 of pointers. A pointer may be enqueued in the FIFO of pointers in block 170 by asserting the "push" control line active and placing the pointer (SRAM 130 address) on the "push-addr" input lines to block 170. Similarly, a pointer value may be dequeued by asserting the "pop" control line on block 170 with the pointer value appearing on the "pop-addr" bus. Block 170 has a "full" output that when active 15 indicates that no more pointers may be enqueued because all the registers 150_{1..p} are occupied and the queue is at its maximum depth, p. Block 170 also has an "empty" output that indicates when no pointers remain enqueued.

In a preferred embodiment, the RAM 130 is divided into two arrays, an address control array AC and a data array D. The address control array AC stores 20 the control information 202 and the address 204 for each entry and the data array D stores the data 206 for each entry. Dividing the non-blocking load buffer into these two arrays AC and D realizes a more efficient chip layout than using a single array for both purposes.

Figure 10 illustrates the entire structure for a non-blocking load buffer in an 25 embodiment of the present invention. Each of the pending queues 114_{0..n} and the return queue 116 may be composed of copies of block 170 illustrated in Figure 8 for example. The queue entries are stored in the AC and D arrays 501 and 502, with pointers to the array locations in the flip-flops that compose the pending queues 114_{0..n}, the return queue 116 and the free pool 112. For the embodiment illustrated

in Figure 10, the write ports to the AC and D arrays are time-division multiplexed using multiplexers 301 and 302 to allow multiple concurrent accesses from the processors $30_{0..m}$, which issue load and store requests over bus 41, and the peripheral devices $50_{0..n}$, which return requested load data over bus 62. Similarly, 5 the read ports are time-division multiplexed by multiplexer 303 between returning load data going back to the processors $30_{0..m}$ via bus 42 and issuing of load and store requests to the peripherals $50_{0..n}$ via bus 61. A processor request control unit 510 accepts the load and store requests from the processors $30_{0..m}$, a return queue control unit 520 controls the data returned by the peripheral devices $50_{0..n}$ to the return 10 queue 116 and a pending queue control unit 530 controls the selection of the pending queue $114_{0..n}$.

When either a non-blocking load buffer or a conventional single FIFO load buffering operation is used, the uncertainty and memory latency is further compounded because individual processors may transmit more than one memory 15 request at a time as illustrated in Figure 3(a) for example. In Figure 3(a), processor A issues three non-blocking loads to a memory device after a first non-blocking load is issued by processor B but before a second non-blocking load is issued by processor B to the same memory device. The first memory transaction issued by processor B, labeled B1, will be executed immediately. However, the latency of 20 the second transaction, labeled B2, by processor B, t_{B2} , will be considerably longer than the latency of the first transaction, t_{B1} , since the second transaction B2 will not be handled by the memory system until all of the requests A1, A2, and A3 by processor A have been processed.

One consideration for overcoming all of these unpredictable memory latency 25 effects is to have the programmer carefully design the times at which each real-time process attempts to access the memory which effectively has software serialize access to the memory banks instead of hardware. For instance, two processors may alternate memory accesses at regular intervals to match the throughput capabilities of the memory system. However, such programming places a heavy burden on both

the program and the programmer because the programmer must be aware of the detailed temporal behavior of the memory access patterns by the program. The behavior of the programs is usually very complex and may change based on its input data, which is not always known. Therefore, it is impractical for the programmer 5 to know the detailed temporal behavior of the memory access patterns by the program.

Also, because the combinations of processors that are making the requests are even more difficult to predict and the combinations of programs running on the different processors might not always be the same on each occasion or even known, 10 the detailed temporal behavior of the memory access patterns by the program may not be known. For example, two processes that run two different programs on two different processors may have been written by different programmers and the source code may not even be available for either or both programs. As a result, this programming approach for overcoming the unpredictable memory latency defeats 15 the purpose of using a non-blocking load buffer, which was designed to simplify the burden on the programmer and the compiler by relaxing the scheduling constraints under which memory operations can be issued while still permitting efficient utilization of the processor.

Figure 9(a) illustrates another embodiment of the present invention for a 20 multiple-priority version of the non-blocking load buffer. In this embodiment, the non-blocking load buffer is a variation of the basic non-blocking load buffer where multiple pending sub-queues $214_0, 214_1, \dots, 214_p$ exist as a component of each pending queue $114_{0..n}$. The outputs of the sub-queues $214_{0..p}$ are then input to the sub-multiplexers 218 associated with one of the peripheral devices. Each output of 25 the sub-multiplexers 218 are then input to the main multiplexer 118. Each of these pending sub-queues is assigned a unique priority level. In the simplest implementation of this multiple-priority version of the non-blocking load buffer, illustrated in Figure 9(b), there are two pending sub-queues 214_0 and 214_1 for a peripheral device with sub-queue 214_0 being assigned a high priority and sub-queue

214₁ assigned a low priority. The multiple priority non-blocking load buffer issues memory or peripheral transactions in a highest-priority-first manner. In other words, no transaction will be issued from a pending sub-queue for a given peripheral unless all of the higher priority sub-queues for that same peripheral are 5 empty. For each peripheral device, requests are issued from the pending sub-queue in a FIFO manner for memory transactions of the same priority.

The computational processes associated with the memory transactions that take place at each of these relative priority levels execute on the processors 30_{0..m}. The priority levels are assigned by the processors based on the scheduling 10 parameters of all the processes in the system. The priority level of the process is designed to reflect the relative importance of the memory accesses with respect to allowing all of the processes to meet their deadlines if possible. The priority levels may also be applicable in a uni-processor environment. For example, interrupt 15 handlers may be provided to achieve low latency by using higher priority memory transactions than the transaction being interrupted. The priority level may be identified by adding a priority tag to the memory transaction. This priority tag is used to channel the memory transaction into the pending sub-queue with the matching priority level, thus the selection of the appropriate destination pending sub-queue for a non-blocking memory transaction is a function of both the address and 20 the priority level of the access. The priority tag may be stored with the other control information 202 in the queue entry 201 corresponding to a given non-blocking memory access.

Figure 11 illustrates a flow chart for the functions performed by the pending 25 queue control unit in an embodiment of the present invention for the multiple-priority version of the non-blocking load buffer. At step S10, the counter i for a pending queue is initialized to zero. At step S20, the counter is compared to the number n of pending queues 114. If the counter is not equal to the total number of pending queues, the counter is incremented at step S30 and a determination is made at step S40 of whether the peripheral corresponding to the counter is ready. If the

5 peripheral device is determined not to be ready at step S40, the process returns to step S20. However, if the peripheral corresponding to the counter is ready at step S40, a pending queue is selected at step S50 with the highest priority that also contains at least one memory transaction. At step S60, memory transactions are issued in a FIFO order from the selected sub-queue and the process returns to step S20. If the counter is equal to the number of pending queues at S20, the process has been completed for all the pending queues and step S10 is returned to where the counter is initialized.

10 Figure 3(b) illustrates memory latency for the multiple-priority non-blocking load buffer. In this example, processor B's memory transactions are assigned a higher priority than processor A's memory transactions. Therefore, transaction B2 is delivered to the memory before transactions A2 and A3 even though the request to begin transaction B2 arrived at the non-blocking load buffer after requests A2 and A3. As a result, the latency for transaction B2, t_{B2} is less in Figure 3(b) than t_{B2} in 15 Figure 3(a), which illustrates a non-blocking load buffer that does not offer the benefit of multiple-priority scheduling. Using the multiple-priority version of the non-blocking load buffer, Processor B spends less time stalled waiting for transaction B2 to complete as illustrated by the comparison in Figures 3(a) and 3(b).

20 These priority levels may be heuristic in nature. For example, if using earliest deadline first (EDF) scheduling, a process should not be assigned a lower priority than the priority of any process which has a more distant deadline. In general, this priority level is not necessarily fixed for each process and the priority level may vary over time as the demands of the real-time process or of other processes change. As another example of a priority selection mechanism, if the 25 load-use distance is known for a load instruction (as computed by a compiler), it can be used to set a priority for each individual memory access instruction. (Higher load-use distances result in lower priorities.) In general, non real-time processes (the processes having no deadlines) are typically given the lowest priority.

Process priority may also be used to arbitrate for the limited resources within the non-blocking load buffer itself. For example, in one embodiment of the non-blocking load buffer, limited data storage memory is shared among all peripheral devices and processors. When the total number of slots among all of the non-blocking load buffer queues exceeds the number of non-blocking memory entries, it is possible for a process at a low priority to prevent a higher priority process from completely utilizing one or more of its pending sub-queues by using up these entries. The use of priority in allocating non-blocking memory entries can be used to eliminate this effect. For example, the maximum number of outstanding non-blocking memory transactions may be specified for each of the available priorities.

Once an appropriate priority has been determined for a process, the priority of its memory transactions might be specified to the non-blocking load buffer by employing any of several techniques. For example, the priority of any given memory transaction might be determined by an operand to the instruction that performs the memory access or the priority can be associated with certain bit combinations in either the virtual address, the physical address of the memory or the physical address of the peripheral device accessed. Alternatively, the processors might be designed with programmable registers that set the priority of all memory accesses for each processor. Yet another possible technique is to store the memory access priority in the page table of a virtual memory system and thus make the memory access priority a characteristic of specific memory pages. When these various techniques are combined with conventional memory management and processor design techniques, memory priorities can be treated as privileged resources. As a result, the operating system reserves the highest priority levels for its own real-time tasks and therefore the user level programs are extremely limited in their ability to disrupt important system real-time tasks.

To maintain memory coherency in a multiple-priority non-blocking load buffer, which is not always necessary, no load or store may be issued to any of the peripheral devices while a load or store to the same address is outstanding or while

an earlier load or store to the same address is pending. The memory coherency may be maintained by time stamps combined with address comparison. However, memory coherency may also be maintained more simply by ensuring that requests of different priorities are sent to non-overlapping address segments within the same 5 peripheral device.

In the normal operation of programmed I/O activity, the processors 30_{0..m} do not need to exactly schedule loads to achieve maximum throughput from the peripheral devices but can instead burst out requests to the limits of the non-blocking load storage and, if properly programmed, perform useful work while 10 waiting for data to return. The non-blocking load buffer allows application code to use programmed I/O (memory mapped I/O) for achieving near maximum throughput from the I/O devices of varying speeds which reduces the effective latency of the I/O peripheral devices and relaxes the scheduling constraints on the programmed I/O. The non-blocking load buffer functions to rate-match the requests from the 15 processor to the peripheral devices and back and to act upon the priority of requests from the processor to allow high priority requests to go ahead of low priority requests already buffered.

The non-blocking load buffer uses queues of pointers to centralize storage to increase storage density, parallel queues to implement requests of different 20 priority and memory segment descriptors to determine priority. Accordingly, I/O requests to a fast device need not wait behind requests to a slow device and requests from several processes of a signal processor running multiple processes are buffered so that the processor is not unnecessarily idled and the time to complete tasks is reduced. Because the multiple priority non-blocking load buffer has multiple 25 pending sub-queues for each of the peripherals, a processor used in combination with this multiple priority non-blocking load buffer is able to run real-time processes of varying deadlines by use of non-FIFO scheduling of memory and peripheral accesses. Also, the multiple priority non-blocking load buffer simplifies the burden on the programmer and the compiler by relaxing the scheduling constraints under

which memory operations can be issued while still permitting efficient utilization of the processor.

The invention being thus described, it will be obvious that the same may be varied in many ways. Such variations are not to be regarded as a departure from the spirit and scope of the invention, and all such modifications as would be obvious to one skilled in the art are intended to be included within the scope of the following claims.

WHAT IS CLAIMED IS:

1. A multiple priority non-blocking load buffer comprising:
 - a plurality of variable depth pending queues for queuing entries of memory or I/O requests generated by processors to peripheral devices, each of said pending queues including a plurality of sub-queues; and
 - 5 a variable length return queuing unit for buffering data returned from said peripheral devices in response to said I/O requests.
2. A multiple priority non-blocking load buffer according to Claim 1, wherein said variable length return queuing unit comprises one variable length return queue.
 - 10
3. A multiple priority non-blocking load buffer according to Claim 1, wherein said variable length return queuing unit comprises a plurality of variable length return queues.
4. A multiple priority non-blocking load buffer according to Claim 3, wherein each of said plurality of variable length return queues corresponds to each of said processors.
 - 15
5. A multiple priority non-blocking load buffer according to Claim 3, wherein each of said plurality of variable length return queues corresponds to each of said peripheral devices.
- 20 6. A multiple priority non-blocking load buffer according to Claim 3, wherein each of said plurality of variable length return queues corresponds to a unique priority level.

7. A multiple priority non-blocking load buffer according to Claim 1, wherein each of said sub-queues within one of said pending queues is assigned a unique priority level.

8. A multiple priority non-blocking load buffer according to Claim 7, 5 wherein said I/O requests include an address and a priority tag, said address directs said memory or I/O requests to a corresponding one of said pending queues and said priority tag channels said memory or I/O requests to a corresponding one of said sub-queues within the one said pending queue.

9. A multiple priority non-blocking load buffer according to Claim 8, 10 wherein each of said pending queues comprises a priority controller for issuing said memory or I/O requests from said sub-queues in a highest priority first manner.

10. A multiple priority non-blocking load buffer comprising:
a variable depth pending queue for queuing entries of I/O requests generated by a processor to peripheral devices, said pending queue including a 15 plurality of sub-queues; and
a variable length return queuing unit for buffering data returned from said peripheral devices in response to said memory or I/O requests.

11. A multiple priority non-blocking load buffer according to Claim 10, wherein said variable length return queuing unit comprises one variable length 20 return queue.

12. A multiple priority non-blocking load buffer according to Claim 10, wherein said variable length return queuing unit comprises a plurality of variable length return queues.

13. A multiple priority non-blocking load buffer according to Claim 12, wherein each of said plurality of variable length return queues corresponds to each of said peripheral devices.

14. A multiple priority non-blocking load buffer according to Claim 12, 5 wherein each of said plurality of variable length return queues corresponds to a unique priority level.

15. A multiple priority non-blocking load buffer according to Claim 10, wherein each of said sub-queues within said pending queues is assigned a unique priority level.

10 16. A multiple priority non-blocking load buffer according to Claim 15, wherein said memory or I/O requests include a priority tag, said priority tag channels said memory or I/O requests to a corresponding one of said sub-queues.

15 17. A multiple priority non-blocking load buffer according to Claim 16, wherein said pending queue comprises a priority controller for issuing said memory or I/O requests from said sub-queues in a highest priority first manner.

18. A method for limiting admission of prioritized memory or I/O transactions for storage in a memory comprising:

(a) dividing the memory into a plurality of queues having unique priority levels;

20 (b) channeling said prioritized memory or I/O transactions to the ones of said queues having its unique priority level corresponding to said prioritized memory or I/O transactions;

(c) storing said prioritized memory or I/O transactions in the corresponding queues when space is available therein; and

5 (d) storing said prioritized memory or I/O transactions in one of said queues having a priority level which is one priority level less than the unique priority level corresponding to said prioritized memory or I/O transactions channeled at said step (b) when the space available in the queues corresponding to said prioritized memory or I/O transactions is exceeded.

10 19. A method according to Claim 18, wherein the memory is contained within a non-blocking load buffer.

20. A method according to Claim 18, wherein said memory or I/O transactions comprise loads and stores.

ABSTRACT OF THE DISCLOSURE

A non-blocking load buffer is provided for use in a high-speed microprocessor and memory system. The non-blocking load buffer interfaces a high-speed processor/cache bus, which connects a processor and a cache to the non-blocking load buffer, with a lower speed peripheral bus, which connects to peripheral devices. The non-blocking load buffer allows data to be retrieved from relatively low bandwidth peripheral devices directly from programmed I/O of the processor at the maximum rate of the peripherals so that the data may be processed and stored without unnecessarily idling the processor. I/O requests from several processors within a multiprocessor may simultaneously be buffered so that a plurality of non-blocking loads may be processed during the latency period of the device. As a result, a continuous maximum throughput from multiple I/O devices by the programmed I/O of the processor is achieved and the time required for completing tasks and processing data may be reduced. Also, a multiple priority non-blocking load buffer is provided for serving a multiprocessor running real-time processes of varying deadlines by prioritization-based scheduling of memory and peripheral accesses.




Fig. 1(a)




Fig. 1(b)




Fig. 2(a)




Fig. 2(b)




Fig. 2(c)




Fig. 3(a)





Fig. 3(b)

Fig. 4






Fig. 5




Fig. 6




Fig. 7(a)




Fig. 7(b)




Fig. 7(c)

Fig. 7(d)

Fig. 7(e)

Fig. 7(f)

Fig. 7(g)




Fig. 8




Fig. 9(a)




Fig. 9(b)




Fig. 10




Fig. 11